



# Core Computing

## (Infrastructure + Software)

# Status

---

Robert Hatcher

Fermilab Computing Division


MINOS Ely Week-in-the-Woods Collaboration Mtg 2011-06-15



# Effort



- FNAL (Arthur K. + ~~Robert H.~~)
  - Now part time on Minerva/NOvA
  - task/skip separation/overlap/holes
- Retirement/Wind-down
  - Nick W., George I./ Sue K., **Robert H.**
- Help from others:
  - Rashid M., Howie R. (batch)
  - Jarek Nowak (simulations)
  - Jon Paley (DatabaseMaintenance)
  - many collaborators ... (now and past)





# Recent Infrastructure Changes



- BlueArc disk

- /minos/scratch → /minos/app (“scratch” no longer anywhere)
- 4 daily rolling snapshots for both app and data disks
  - /minos/data/.snapshot/2011-06-11\_1701-0500.Minos-Data-Daily
- 30TB more due to be allocated to MINOS this year
  - don't use it too quickly please; clean up old kruft

- MINOS DB servers saga

- minos-mysql2 restored as master db (freeing up minos54)
  - failed Jan 20th, restored May 19th
- minos-mysql1 (farm db) failed June 9th
  - was old and very underpowered for the task
  - keep-up reco processing (dogwood4) temporarily using master
- configuring minos54 as replica db (as minos-farmdb)
  - use MySQL replication instead of DBAuto for sync w/ master
  - ensuing discussion on “frozen db” (e.g. dogwood1 priming set) vs. using DBI rollback mechanism to accomplish the same
- also shared for Minerva + Argonne MINOS MC generation

Robert Hatcher





# Recent Infrastructure Changes



- Miscellany
  - CRL → ECL, memopad → redmine wiki
    - generally accessible via “Services” account (pswd != kerberos)
    - ECL non-std accounts copied from CRL w/ generic password
    - CRL is frozen (content was copied to ECL); to be removed in July
  - increase in grid slots
  - ongoing question about ganglia support
    - currently split to separate web areas
  - CPN
    - lock file snafu (MINOS saw in late May) won't re-occur
    - 6.7 locks served to MINOS



# Imminent Infrastructure Changes



- minos50/51 move to SLF5 tomorrow
  - previously: GPFarm grid moved to SLF5; also minos52/53
  - existing code built on SLF4 should still run everywhere
  - base release builds will (generally) continue on minos27 (SLF4) for a while; test releases on SLF5 should be compatible
  - reference virtual machine: minos-slf4 (also slf5)
- nightly “development” default: gcc 4.5.1
  - gcc 3.4.3 will continue to be built
  - use: `minos_setup -q GCC_3_4`
  - rebuild test releases or they’ll appear to have “disappeared”
    - reminder: regularly rebuild if base release is “development”



# Imminent Infrastructure Changes



- Condor/Grid

- upgrade to condor 7.4 tomorrow
  - worker nodes already done; master node tomorrow
  - allow more jobs; no need for MINOS throttling; cluster job [rate] limiting
- move to common “jobsub” - maintain by CD/REX for all v expt.
  - staged adiabatic transition; volunteers needed
  - access to ~100 gpcf slots (similar to MINOS cluster, batch only)
- CD/REX support
  - really, this time we mean it

- Miscellany

- minos26 to be retired soon
- Control Room
  - system maintenance assumed by CD/FEF (was Jon Urish in REX)
  - software support for remote shifts being finalized
    - local copy to be installed to replace AFS?



# Near Future Infrastructure Changes



- FermiMail IMAP
  - Exchange 2010 w/ proper IMAP (not Exchange 2007)
  - staged migration July through October
- Guest wireless network
  - toxic pool; details still unknown
- SAM Projects?
  - 21 projects since Nov 2010; all Nathaniel's LinearityFit
  - "projects" require infrastructure beyond just basic SAM
  - Arthur has a prototype alternative
- Cleanup of condor-tmp & condor-exec
  - one user has 400K files in a single directory
  - reminder e-mails aren't cutting it
    - one simple solution is to push to subdirs YYYY-MM





# Near Future Infrastructure Changes



- FermiMail IMAP
  - Exchange 2010 w/ proper IMAP (not Exchange 2007)
  - staged migration July through October
- Guest wireless network
  - toxic pool - no system scan/block; details still unknown
- MINOS-CD MOU
  - version being drafted, reformatted w/ emphasis change
- File ownership by owner rather than minosana
- Group accounts
  - mindata, minsoft, minosraw, minospro, minosana
- dCache/Enstore
  - pin raw data to dcache disks (avoid tape access)
  - want 120TB disk for small files w/ enstore access but not tape





# Software

- Package Maintenance
  - people seem to be responding to the e-mail: **thanks**
    - exit interviews for those leaving the collaboration?
- not much new development on core packages
- New tagged releases to use GCC 4.5.1
  - R2.6 is gcc 4.5.1 only (but built on SLF4)
  - S1 1-04-29-R2-06 is gcc 4.5.1 only, on SLF5
  - both gcc 3.4.3/4.5.1 for nightly development
  - swapping default for setup\_minos w/o “-q”
    - -q GCC\_3\_4 (implied) // -q GCC\_4\_5 (alternative)



# Software Issues

- N00017725\_0001.mdaq.root
  - first seen to fail in (gcc4.5+SLF5) test in snapshot release S11-04-29-R2-06
  - also fails for various combinations of SLF4, gcc 3.4.3, even OS X 10.6 (Mac Snow Leopard)
  - depends on prior processing history; skipping events changes behaviour (~3800 records in)
  - fingers point to lots of places:
    - AlgFitTrackCam (CleanNDList)
    - AlgRmMu
    - Persistency/ROOT
    - Calibrator



# Robert Tasks

- filled BEAMMONSPILL DB table (POTs)
  - missing spills due to firewall drop-out (recovered on Feb 8th):
    - 2010-08-22 02:17:11 to 2010-09-17 16:09:29 (Minerva special runs)
  - high packet load/“storms” on beam network (recovered on Mar 4th)
    - 2011-02-10 22:00:00 to 2011-02-11 02:59:59
    - 2011-02-17 06:00:01 to 2011-02-23 05:59:58
  - hardware failure timeout (recovered on Apr 26th)
    - 2011-04-20 05:00:00 to 2011-04-26 04:59:53
  - ACNet data extracted by Phil Adamson
  - task force for NuMI beam line daq (other v expt)
- Alignment studies: Jasmine
  - generation done for 5 rotations schemes
    - detector not built perfect
    - our knowledge of the construction not perfect
  - my effort: ~~helping, not~~ doing generation





# Other Robert Tasks (2)

- EventKinematics (Hops MC, prototype for NOvA et. al.)
  - have ability to select non-physical fiducial volume
  - neutrino flavor changing framework (missing from GENIE)
- Detector Modeling
  - rock model (cavern+tunnel is not really a box)
  - new material upstream of MINOS (e.g. Minerva)
    - outstanding request for simple “soup” of material
    - still no values...
    - complications when Minerva overlays




# Robert Future

- Internal CD Migration to Simulations Group
  - I'm not disappearing off the face of the earth
  - group want me to remain plugged in to v needs
- my “Operations” effort: helping; **not** doing!
  - nightly builds, release tags + builds: need volunteer
    - writing up a note/wiki
  - no more running jobs (misalign MC sufficient?)
  - limited consulting on framework/reco maintenance
  - users should stop beating on database
    - hopefully we now understand triggers causing issues



# The Hall

- MINOS Near Hall model is not very accurate
  - just a rectangular box
  - no Minerva!







# Extra Stuff

- Here be dragons...



# Transitions

- Art is participating in Near Detector Transition Task Force (MINOS + Minerva)
  - software component as well as hardware
    - data handling, calibration, data quality, batch processing, simulation, code maintenance, ...
- Remote Shift software
  - not officially supported by FNAL-CD (at this time)
  - numerous issues to still be worked out fully
    - especially security issues
  - would like to see it all packaged in an easily deployable bundle



# R2.0.5 - Unified Dogwood

- Desire to use same code for “keep-up” and “physics” reconstruction
- Dogwood4 = “keep-up”: R2.0.5 + dogwood0 DB
  - use old calibrations to generate new calibration
  - no one should use it for physics
  - better than Dogwood0 (R2.0.0)
    - uses code that got fixed for Dogwood3
      - SEGVs, CalDrift f[Front]Median, muon range bias + sign flip
    - fewer crashes  $\Rightarrow$  happier batch group
- Dogwood5 = “physics”: R2.0.5 + dogwood1 DB
  - more robust/faster code than Dogwood3 (R2.0.3)
  - correct filling of Nathaniel’s timing block





# R2.0.0 $\Rightarrow$ R2.0.3 $\Rightarrow$ R2.0.5

- R2.0.0 to R2.0.3 (see previous slide)
  - fix 100%/pass FarDet cosmic failures in Dogwood0
- Reco Modules fetch right RawData Record
  - need RawDataDigit? ask for DaqSnarl record explicitly
  - allows Persistency to use kLowerBound to keep appropriate DaqMonitor record in MOM
- Sue K.'s speedups
  - CandNtupleSR/NtpSRBleachFiller; CandFitTrackSA/MatrixCalculator
- Actual bugs fixed
  - AlgAltDeMuxBase
    - caused 100% Dogwood3 failure for MC on SLF5 (Andy B.)
  - DbValidityRecBuilder (backport from development, Nick D.)